Thesis/
Reports
Bishop,
I. D.
c. 1

SIMULATION AND ANIMATION OF RIVERS AND STREAMS: A FEASIBILITY STUDY

by
I.D. Bishop, E. Flaherty and K. Aldrich

12/31/90

SIMULATION AND ANIMATION OF RIVERS AND STREAMS: A FEASIBILITY STUDY

Cooperative Agreement No. 28-CO-533

between

TEXAS A&M UNIVERSITY, COLLEGE OF ARCHITECTURE

and

ROCKY MOUNTAIN FOREST AND RANGE EXPERIMENT STATION
Research Work Unit No 4851

FINAL REPORT: DECEMBER 1990

prepared by Ian D. Bishop, Eric Flaherty and Kipp Aldrich, Visualization Laboratory, College of Architecture, Texas A&M University

This research explores the feasibility of alternative approaches to the simulation of moving fresh water bodies in complex terrain. The candidate approaches identified initially were: mapping, dynamic modeling and live image merging. This report gives details of findings within each of these categories. The accompanying videotape serves as illustrative material for the report. The report should be read in conjunction with the videotape.

Method 1. Mapping

Several rivers in Texas and Colorado were videotaped from a number of different angles. For this first study a section of the Brazos River, Texas captured from directly above was chosen because the high contrast gave a very clear pattern of movement (see Tape sequence 1)* . The camera/recorder was a SONY Broadcast BVP-7/BVV-5 BetaCam-SP with 550 lines of horizontal resolution in component mode.

Thirty frames (1 second) of motion was transferred from the tape to a digital frame store and then to computer disk. Each of the thirty frames was converted to a color map and applied as a texture sequentially to a hypothetical river surface. Each frame took about 3 minutes to render on a Silicon Graphics 2D/240 dedicated computer. Because the start and end frames were not identical a distinct discontinuity occurs each 30 frames (i.e. each second) in the rendered sequence (Tape sequence 2). A number of procedures were tested in an attempt to overcome this problem - these are described in the Assessment section below.

In Tape sequence 3, the 'camera' is moving across the water surface. This takes attention away from the annoying 'pop' where the texture maps cycle back to the first frame. This seems to be a consequence of the additional visual information to be absorbed. The observer is concentrating on his/her own 'movement' rather than that on the river which becomes more a secondary impression that the major point of attention.

An alternative to use of a color map is to translate the color variation of the images into variations in reflectance angle, i.e. a bump map, and apply this to the river surface. Under

---

* Tape sequences 1,2,4,5,6,7 and 8 are shown first at full speed (30 frames per second) and then repeated at 15 frames per second.

LIBRARY COPY
ROCKY MTN. FOREST & RANGE
EXPERIMENT STATION

high illumination the river acquires similar lightness/darkness patterns as the color map but derives its color more from the sky and surrounding landscape that from the original imagery. Note in Tape sequence 4 the reflection of the land at the water's edge. The water in this case is behaving as an optical entity with appropriate optical characteristics (e.g. a refractive index of 1.5). The option took about 22 minutes per frame to render: considerably more than the color map approach.

Combining both the color and bump maps into a single model (sequence 5) gave some additional advantages (e.g. a greater sense of relief in the water surface).

Each of these maps used the original imagery undistorted. This meant any apparent response of the water movement to the terrain was purely co-incidental (e.g. the water did not necessarily accelerate at corners or through narrows). It was then found that if the model of the water surface was distorted after the map application that the map distorted also. This meant, for example, that where the polygons of the river surface were stretched parallel to the flow direction the water appeared to move faster. As Tape sequences 6 and 7 show this has the potential to allow a proper fit of water movement characteristics to terrain characteristics. It is however rather time consuming for an operator to generate the proper distortions using the current modeling interface. With more operator practice, or a grid distortion algorithm based on computed flow characteristic (see below), better results in terms of both time taken and quality of simulation can be anticipated.

The difference between sequences 6 and 7 is that in 7 the map covers the full width of the stream and the tiling problems apparent in 6 disappear. This presumes however that a map covering water area corresponding to the dimensions of the hypothetical river is available.

In Tape sequence 8 all the most effective techniques in the previous sequences are combined. The images show a combination of color and bump mapping and texture distortion based on full width coverage. The effectiveness of the texture distortion process is seen most effectively on the left side (as we see it) of the island and on the right side of the stream behind the island. In both cases the appearance of accelerated water motion is achieved and coincides with the expectation of faster flow due to the narrower stream bed in the first case and stream bending in the second. Problems remain however. For example, on the left near corner of the island water appears to be accelerating from out of the rock face.

One of the most serious problems with this approach is the visual 'pop' at the transition from the last frame of the captured sequence back to the first. To overcome this a number of options were explored. First, the sequence was reversed in alternate seconds but this made the water appear to wash back and forth rather than flow in a steady direction. Second, the source video was searched to find alternative end frames more like the start frame. Although some improvement is possible the water is never quite the same twice and some 'jump' will always be apparent. Third, a routine was written to interpolate between the last and first frames, but this approach was also abandoned as the water seemed to stop moving and lose contrast during the change over. More complex transition algorithms might eventually be possible but this would be a substantial development effort. Finally, the option exists to capture a sequence as long as the required simulation. The problem here is that with finite available disk space only a few seconds could be rendered at a time, files would then have to be cleared off the computer system and another few seconds of imagery generated. As one of the objectives of this approach was to limit the amount of operator intervention in the process this last option was also unattractive.

The use of bump maps is appealing because it makes incorporation of reflection effects possible - note the reflection of the land form in the water in sequences 5 and 8. This does however increase rendering times significantly. Each frame in sequence 8 took about 3 hours to render.

Method 2. Dynamic Modeling

This procedure, which is necessary for precise simulation of continuously varying river flow in its normal environment, requires hydrological modeling. River flow models come in varying degrees of sophistication. One-dimensional models predict only an average velocity and water surface elevation across each river cross-section. They do not respond well to changes in channel dimensions, roughness, obstructions or other complex conditions common in the natural world. Two-dimensional models provide a more accurate determination of river levels and velocities. For rivers, a depth averaged model with flow in the horizontal plane is commonly used. A copy of the finite element based two-dimensional model developed by the USGS for the Federal Highway Administration (Froehlich, 1988) has been acquired for evaluation in the context of this project. This software is known as FESWMS-2DH and runs under MS-DOS.

Based on ground surface contours and total flow volume FESWMS-2DH generates surface velocities and levels at each node on a finite element net. These values are, however, averaged over time and no transient surface movement characteristics are generated.

To precisely generate transient surface phenomena - turbulence, waves etc - a complete three-dimensional turbulence model is required. The ASCE Task Committee on Turbulence Models in Hydraulic Computations (1988) reviewed the computer resource requirements for implementation of such models. The computation time per cell per iteration depends on the equations used and the speed of the computer used. The minimum in any of the models reviewed was about 0.5 milliseconds on a CDC 7600 and an IBM 3033. According to Dongarra (1989) these machines have MFLOPS ratings of 2.0 and 1.7 respectively. A single processor CRAY Y-MP/832 has a rating of 84 MFLOPS and so could be expected to reduce the iteration time to about 0.01 milliseconds. To visually simulate surface water turbulence a cell size of no more than 1 cm is probably necessary. Thus, to model an area of 2m by 2m to a depth of 0.25 m and generate a new surface 10 times per second for 10 seconds would involve 100 million cell-iterations or about 17 minutes of CRAY time.

Such a detailed polygonal surface would also have quite long rendering times and this approach appears to be unrealistic for fully detailed surface rendering. There are however intermediate stages in the use of dynamic modeling which could be pursued.

For example, once an averaged surface and velocity have been generated using a two-dimensional model these could be used as input parameters to one of the techniques being employed by computer animators to produce water surface simulations. Perlin (1985) used normal perturbations generated by spherical wave fronts arising from random point source centers. This allowed a watery surface to be placed on any object. Mastin et al (1987) describe a Fourier transform based approach to ocean surface synthesis. Peachey (1986) developed a model suitable for rendering and animation of wave approaching and breaking on a sloping beach. He used a height field computed from the Stokes wave model (1957) and included wave refraction as the water depth decreases. Fournier and Reeves (1986) also created algorithms for ocean surface simulation with breaking waves and foam. Their basic model was the Rankine wave (1863) where particles of water describe circular or elliptical orbits. Their disturbing force was the wind and restoring force gravity. Various wave shapes can be produced by changing the parameter of the orbits.

These techniques produce very effective images of open ocean and shore lines. They may also be effective on large bodies on inland water, including major rivers but would have similar edge problems to Method 1.

Tape sequences 9,10 and 11 show implementation of the Mastin et al (1987) Fourier transform approach. White noise is created over a two dimensional grid. This is Fourier transformed to give a white noise distribution of frequencies. A frequency filter based on the distribution of frequencies observed on the ocean surface is applied. An inverse Fourier transform is applied to the filtered frequency set and the result is an undulating surface similar to a wind driven water surface. The is seen under the Method 2 title. By introducing a phase transition during the first transform the peak and troughs move up and down and a sense of water motion is created (sequence 9). It was felt however that this animation did not have the sense of flow direction we were seeking for river conditions. This first sequence had no phase shift/positional shift in the inverse transformation. In the second sequence (#10) a constant factor was added in both X and Y directions during the inverse transform This gave a greater sense of motion although the use of a smaller dimensioned grid reduced the image quality. In the ideal situation the apparent motion of the water surface (and the amplitude of the waves) should be controlled by output from a dynamic model such as the Froehlich (1988) model described above. In order to assess the potential for such linkage a variable shift was introduced in the inverse transform. This simply made the X shift dependent on the Y position. The result (sequence 11) is difficult to interpret but does suggest that with further experimentation, special surface effects may be achieved by using variable shift factors. Once the relation between the magnitude of these factors and their visual effect is better understood it would be a relatively simple matter to define a matrix of indices which would define patterns across the surface of the water.

Linkage of the index matrix to dynamic stream models could then follow although there would still be problems to overcome. These would include the difficulty of executing Fourier transform based procedures on stream grids which are variable in length and width and could include islands and other features having distinctive effects on the water surface.

Method 3. Image merging

The final approach relies on merging of existing video sequences of actual rivers. Conceptually, a river from one location can be merged with the background from another location. The particular application envisaged for this project however sought to change flow levels in a single river without the changes in background (leaf conditions, lighting, sky conditions etc) that exist on different days and in different seasons.

Sequences 12 & 13 show the results of application of this technique at two sites on the Cache La Poudre River in Colorado. The original videotapes were collected from 22 sites at 7 different flow levels by the Forest Service for an earlier project. Sequence 12 is the upstream view from site 7 at 2303 cfs, 1202 cfs and 152 cfs, while sequence 13 is the upstream view from site 12 at 152 cfs and 2303 cfs. Essentially, the process involved:

(a) capture of a still image from the sequence with the highest flow level (a 35mm slide could be scanned to provide the background if the lens configuration gave an identical field of view to the video camera);
(b) manual drawing of a mask covering the background portion of the scene;
(c) blurring of the edge of the mask;
(d) running the video of a lower river level behind the captured and masked background; and
(e) recording the new composite image on videotape.

Sound has not been included at this stage but can be added from the original source tapes to match the different river levels.

The simplicity of this description hides a multitude of technical problems - all of which, as can be seen, can be overcome. Some of these arise from the original imagery and some are technical issues that result from working in the video medium. The original imagery should all be taken in the same location, with the same tilt angles on the camera and the same lens zoom factor, i.e. there should be some precise registration mechanism. It is best if lighting conditions are similar because the illumination on the water and exposed banks must match, to a fair degree, the background. In cases where these equalities do not all exist (as in the examples shown) they can be corrected electronically but this increases the work involved considerably.

The technical capabilities necessary to achieve the best results include:

- linear keying (simulated using the blurred mask edge)
- a properly timed and referenced video production system
- source tapes of the highest possible resolution and signal-to-noise ratio (i.e. original tapes in a high quality video format)
- a broadcast quality frame grab/frame buffer with appropriate interface electronics to capture video from tape
- a frame buffer with alpha channel (i.e. information bits beyond the color data bits) output for masking
- appropriate paint/image processing software to prepare backgrounds stills and masks
- the timing delays induced by external equipment (e.g. encoders) in color signals from the framebuffer must be matched by equivalent delays in the mask signal from the alpha channel at the video production switcher (it was necessary in our case to introduce 230 m of cable between the framebuffer and the switcher because of encoder delays).

References

ASCE Task Committee on Turbulence Models in Hydraulic Computations (1988) Turbulence modeling of surface water flow and transport: Part IV, J. of Hydraulic Engineering, 114,1034-1051.

Dongarra J.J. (1989) Performance of various computers using standard linear equations software in a Fortran environment. Argonne National Laboratory Tech.Memo.23, Argonne, Illinois.

Fournier A. and Reeves W.T. (1986) A simple model for ocean waves. Proceedings ACM Siggraph 20,75-84.

Froehlich D.C. (1988) FESWMS-2DH Finite element surface-water modeling system: two-dimensional flow in a horizontal plane. Users manual. U.S. Department of Transportation, Publication FHWAS-RD-88-177.

Mastin G.A., Watterberg, P.A. and Mareda J.F. (1987) Fourier synthesis of ocean scenes. IEEE Computer Graphics and Applications, March '87,16-23.

Peachey D.R. (1986) Modeling waves and surf. Proceedings ACM Siggraph 20, 65-74.

Perlin K. (1985) An image synthesizer. Proceedings ACM Siggraph 19, 287-296.

Rankine W.J.W. (1863) On the exact form of waves near the surface of deep water. Phil.Trans.Roy.Soc. A 153,127-138.

Stokes J. (1957) Water waves: the mathematical theory with applications. Interscience, New York.